

Laboratory Automation: The Design Philosophy of the NIST Frequency Measurement Service

Michael A. Lombardi
Time and Frequency Division
National Institute of Standards and Technology
Boulder, Colorado

ABSTRACT

The microprocessor has changed the way that many devices are built and used. Microprocessors have made these devices work better by automating tasks that were formerly done manually. It has also made these devices easier to use. Since the device now does more, the user has to do less.

Microprocessors are now commonly found in consumer products like cameras, stereo systems, televisions, and automobiles. They do things automatically that you used to have to do yourself, and make these products easier to use. This concept now extends to the calibration laboratory, with the advent of the automated calibration system. These systems have many benefits. They are easier to use and learn, and increase the lab's productivity.

In keeping with this trend, NIST has offered an automatic frequency measurement system to the public since 1984. Since that time, NIST has redesigned and improved the system, in a continuing effort to make frequency calibrations easier to perform and understand. This paper discusses the design philosophy behind the system, what it does, how it works, and some ways the system can be enhanced in the future.

INTRODUCTION

Automation brings technology to more people by shifting the burden from the user to the device. In the world of electronics, most automation can be credited to the use of the microprocessor, which allows small computer systems to be embedded just about anywhere. As a result of automation, tasks that used to require a lot of time and effort (like frequency calibrations) can now be easily accomplished.

When NIST decided to design an automated system, their primary goal was making frequency calibrations easier. No one had automated the frequency calibration process. Several manufacturers already offered the pieces needed to calibrate oscillators, but they had failed to offer an automated system. NIST offered the automated system. NIST used

commercially-available hardware to build the system, and then wrote the software needed to tie everything together. The result was an easy-to-use system that any lab could use to routinely calibrate oscillators.

HOW IT WORKS

Like all calibrations, a frequency calibration is simply a comparison. You compare the device you want to calibrate to a *reference* whose performance is known. The device being calibrated is usually a quartz, rubidium, or cesium oscillator. The reference is usually another oscillator (of higher performance) or a frequency signal received by radio. In the case of the NIST system, the device being calibrated can be any oscillator. The reference is a radio signal from the Loran-C navigation system.

Both signals are plugged into a device called a *time interval counter* (TIC). A TIC measures the time interval between the two signals. One signal serves as a start pulse to the TIC, and the other serves as a stop pulse. The TIC starts measuring the time interval when the start pulse arrives, and stops measuring when the stop pulse arrives.

The readings obtained with the TIC need to be recorded. The NIST system records them automatically using a computer. The TIC is interfaced to a personal computer that reads the time interval count from the TIC, and displays and records the number. A series of 5 time interval counts might look like this:

45.67
45.68
45.69
45.70
45.71

Each number represents one time interval measurement. The unit is microseconds (millionths of a second). For example, at the time of the first measurement the two signals were 45.67 microseconds apart. The 45.67 has no particular significance. It's just an arbitrary number used as an example. What is significant is the way the numbers change from measurement to measurement. For example, the difference between the first number in the series (45.67) and the last number (45.71) is 0.04 microseconds. This difference means that one signal is moving or drifting relative to the other signal. The amount of this drift is the quantity that a frequency measurement system attempts to measure.

The NIST system measures this quantity by making measurements for a given length of time and adding together the difference between each pair of measurements that it makes

during that period. T is the designation for the time period during which measurements are made. The sum of the differences is designated delta-T (change in time). The equation for relative frequency (RF) is:

$$RF = \frac{\text{delta-T}}{T}$$

To illustrate, let's say that the oscillator being calibrated drifts just 1 microsecond over a 24-hour period. The 1 microsecond represents the change in time (delta-T). The measurement interval (T) is 24 hours, but needs to be converted to the same unit as delta-T (microseconds). There are more than 86 billion microseconds in 24 hours, so the equation becomes:

$$RF = \frac{1}{86,400,000,000}$$

Relative frequency is a measure of oscillator performance. It represents the frequency error of the oscillator relative to the reference. As you can see, as the amount of drift gets smaller, the number used to state relative frequency gets smaller. In other words, the smaller the relative frequency number, the better the oscillator. In this case, the relative frequency is very small and can be expressed as a percentage:

$$.0000000000116\%$$

Since the number is so small, it is usually converted to scientific notation:

$$1.16 \times 10^{-11}$$

Or, displayed like this on a computer or calculator:

$$1.16E-11$$

This number is used by lab personnel to state the performance of an oscillator. If you hear someone say that an oscillator is off by "one part in ten to the eleventh" this is what they mean. Every frequency calibration system should express this number in some fashion.

THE SYSTEM HARDWARE

We just described how a frequency measurement system works, and mentioned some of the hardware needed to design such a system. This section discusses the hardware used by the NIST system; including the Loran-C receiver, the time interval counter, and the computer. The next section discusses the system software. A block diagram of the entire system is shown in figure 1.

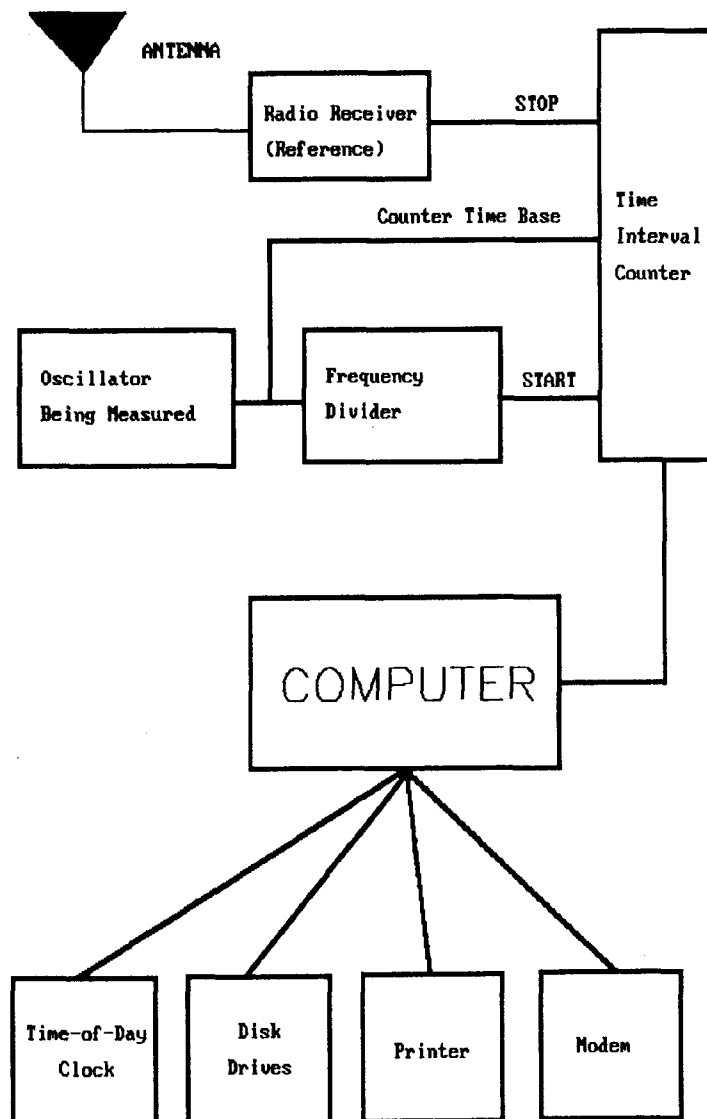


Figure 1 - Block Diagram of NIST Frequency Measurement and Analysis System

LORAN-C RADIO RECEIVER

The NIST system uses a Loran-C navigation receiver for its reference frequency source. Loran-C is a radio navigation system operated by the United States Coast Guard and used by boats and planes to navigate. Loran-C signals are stable and precise, and are an excellent reference for frequency calibrations.

NIST has used several different Loran-C receivers in the system. All of them are commercially available. These receivers are designed for navigation only, and must be modified by NIST before being used for calibrations. However, they work well and are inexpensive, since they are mass produced for the marine market. Some models cost less than \$500.

Loran-C receivers used to be complicated devices to learn and use. However, the usability of recent models has been greatly improved by the addition of a microprocessor. To illustrate this, consider that all Loran-C stations broadcast on the same frequency (100 kHz). The

receiver has a complex task: it must distinguish between signals from many different stations and then pick out the station that you want to receive. It then must locate and track a specific cycle in the Loran-C pulse (usually the third cycle). It used to be a time consuming process to select and track a station. The older receivers were manually operated. Users had to know how Loran-C worked, and had to be skillful enough to adjust the knobs and switches and select the desired station. They then had to attach an oscilloscope or chart recorder to the receiver to identify the correct cycle of the pulse.

Things have changed. Most modern receivers (like the one used by NIST) do all this automatically, through the use of an internal microprocessor and software. Users no longer need to know how Loran-C works, and don't have to spend time studying the receiver manual. Instead, they just turn the receiver on, punch in a number, and the receiver finds the correct station and cycle in just a few seconds. And as an added bonus, the new receivers are smaller and less expensive than the old units. These low-cost, automatic receivers tie in well with the design philosophy of the NIST system.

TIME INTERVAL COUNTER

A time interval counter (TIC) is a fairly common piece of equipment. Several companies sell devices called "universal counters" that perform many functions, including measuring time interval. However, NIST uses a simplified TIC that was designed specifically for use in a frequency measurement system.

The NIST counter has several advantages over other counters. First, it contains some divider circuitry that lets it work automatically with 1-, 5-, or 10-MHz input signals. With an ordinary TIC, you would also need an external frequency divider so that the two signals being compared would be of approximately the same frequency. NIST put the dividers and counter in the same box. Second, while most counters have only one set of start/stop connectors, the NIST counter has four inputs so that up to four different oscillators can be calibrated at once. And finally, the sole purpose of the NIST counter is to measure time interval. It performs no other functions, has no buttons to push or knobs to turn, no panel displays or meters, and is controlled entirely by software. This streamlined design makes the NIST counter an inexpensive and easy-to-use device.

COMPUTER SYSTEM

The current NIST system is controlled by an IBM-PC™-compatible* computer running the MS-DOS™ operating system. Earlier versions of the system used an Apple II™* computer. NIST's design philosophy has been to use an industry-standard personal computer, instead of a dedicated instrument controller. NIST has found many advantages to this philosophy since the service was initiated in 1984, including:

LOWER COSTS. This is due to the economy of scale. Since millions of personal computers are sold each year, they cost

far less than dedicated instrument controllers, which are a relatively small-volume item.

INDUSTRY-STANDARD COMPUTERS HAVE A LONG PRODUCT LIFE. To run this type of service, NIST must be able to buy new computers that are compatible with existing units. If NIST were to use a model that is discontinued and can no longer be purchased, the service would have to stop expanding until the software was rehosted to another computer platform. It's much safer to use an industry-standard model. For example, the Apple II™* computer was first introduced in 1976 and an enhanced version is still available in 1991! The original IBM-PC™* first appeared in 1981. That platform is still going strong 10 years later, and should continue to do so for years to come.

Any PC-compatible (even a laptop) can control the NIST system as long as it has some expansion capabilities. The computers currently used by NIST are standard, off-the-shelf models that are not modified in any way. They are powered by either the 286 or the 386SX microprocessor, and have a color monitor with a VGA graphics card. They have one floppy drive, one hard drive, a tape backup unit, and a dot-matrix printer that prints reports and graphs. Each system also contains a 2400-baud internal modem that lets NIST access the system by telephone, and an interface card for the time interval counter. The computer is housed in a "slimline" case that allows it to fit inside an equipment rack.

THE SYSTEM SOFTWARE

The system software works on two levels. On one level it automates the measurements by getting readings from the counter, storing them on disk, and performing all the functions needed to make calibrations. These functions are basically transparent to the users. They just turn the system on, and the measurements begin.

The second level of software separates the NIST system from other calibration systems. The software uses color displays and graphic printouts to present the calibration results. The user receives data that is more complete and easier to understand than the data produced by other systems.

To illustrate this, remember that a calibration system needs to provide a number for oscillator performance (called relative frequency). This number should be easy for the user to obtain. Some systems present this number on a small LCD or LED display. Rather than present the number in one of the awkward looking formats listed earlier, the designers of the NIST system wanted a way to *visually* represent this number. They

wanted to design a display that users could quickly glance at and get the information they need. This display would greatly simplify using the system, and would be similar to the color-coded meters on the vacuum tube testers that were once found in drug and hardware stores. You just plugged the tube into the socket, glanced at the meter, and quickly found out the condition of the tube.

The display that NIST designed is a bargraph that shows a color-coded bar for each oscillator being calibrated. A picture of this display is shown in figure 2.




Figure 2 - A bar graph showing oscillator performance

The length of each bar indicates oscillator performance (relative frequency). As stated earlier, the smaller the relative frequency, the better the oscillator. For the purposes of the bar graph, the smaller the relative frequency, the longer the bar. In other words, a long bar means a good oscillator. The top of the graph is labeled from 6 through 12. If the bar extends to the 8, the relative frequency is 1×10^{-8} , if it extends to the 9, the relative frequency is 1×10^{-9} , and so on. Obviously, the bars won't instantly show the results, since it might take 24 hours or more to measure the performance of a precision oscillator. But after the system has run for a day or so, the user can check each oscillator's performance with just a quick glance.

The oscillator names are below the bar graph. Each name corresponds to a bar of the same color. The counter readings (and differences) are also displayed. If something goes wrong (a signal is lost, for example) error messages are printed on screen in color to alert the user.

The bar graph approach works well, but NIST realized that they also must provide more detailed graphs of oscillator performance. For many years, calibration labs have used strip-chart recordings (figure 3) to show this information.




Figure 3 - A reproduction of a strip chart recording

The width of the strip chart represents a fixed number of microseconds. For example, the width of the above chart might be 50 microseconds. When an oscillator drifts more than 50 microseconds, the chart overflows. The above chart shows four overflows, or 200 microseconds of oscillator drift. This is the "delta-T" quantity discussed earlier.

To determine the length of the time period when the measurement were made (T), the user must measure the chart length. A typical chart might move at a rate of 1 inch per hour. The user can then calculate the relative frequency using the mathematics discussed earlier. The strip chart does not provide this information. In fact, it does not contain any labels or annotation at all. Everything is left up to the user.

Strip charts are still useful, but they are basically a relic from the pre-computer age. Through software, the NIST system produces much more detailed graphs that let users interpret the data more quickly. Unlike a strip chart, these graphs do not have a fixed chart width. Regardless of how much the oscillator drifts, 10 microseconds, 10 milliseconds, or even 10 seconds, the graph does not show any overflows. Instead, the graph is automatically scaled so that the full range of oscillator drift is shown.

The graphs are also labeled and annotated. The title shows the name of the oscillator being calibrated, and the name of the reference frequency source. The graph shows the amount of oscillator drift (in microseconds). It also shows the start and times for the measurements, the number of the counter channel where the oscillator was connected, and

the relative frequency of the oscillator. This information is all recorded and printed automatically. And as an added benefit, each graph is printed on a full-size 8.5 x 11 inch sheet, so that it can easily be placed in a notebook or file and kept for future reference. A sample graph is shown in figure 4.




Figure 4 - A sample oscillator performance graph

Graphs like the one above are produced automatically every 24 hours. A graph is printed for each oscillator being calibrated. The system also lets users print graphs "on demand" by simply pressing a key on the keyboard. The "automatic" graphs and those printed on demand show the last 24 hours of oscillator data. However, the software allows creating graphs containing from 2 hours to well over 100 days worth of data. All data recorded by the system is stored on the hard disk and then automatically copied to a tape cartridge for backup. It can be retrieved and graphed at any time with just a few keystrokes.

The software provides another graph that shows the "history" of an oscillator. This graph condenses the data and shows just the daily relative frequency values. Each day is indicated by a rectangular marker. If the marker is in the top half of the graph, the oscillator was high in frequency on that day. If it is in the bottom half of the graph, the oscillator was low in frequency. If the marker is on the center line, the relative frequency

was 1×10^{-12} or better. One graph summarizes an oscillator's performance for one month. A sample history graph is shown in figure 5.




Figure 5 - An oscillator history graph

As you can see from these examples, the NIST system presents the user with clear and useful calibration data. In the next section, we'll look at two ways that a system like this could be repackaged in the future.

FUTURE ADVANCES

The NIST system is composed of several instruments placed inside an equipment cabinet, as pictured on the next page. The entire unit is rather large; about 55" inches high, 21" wide, and 17" deep. There are potentially more attractive ways to package this type of system. In this section, we'll look at two alternatives.

1) COMBINING EVERYTHING INTO A SINGLE INSTRUMENT - This requires putting all the hardware into a single box, either a bench type instrument or a rack-mounted unit. The computer would be "hidden" inside the box with the software stored in ROM. It could have a small screen on the front panel, and a small top-mounted printer. The

keyboard could be reduced to a keypad of special function keys, and placed on the front panel. The front panel would also contain connectors for the oscillators. The unit could have an optional Loran-C module for use as a reference. If Loran-C is not included, the unit could still make oscillator-to-oscillator comparisons, and could even include a built-in oscillator.

2) MAKING THE SYSTEM AN "ADD-ON" TO A PC - This system would consist of a time interval counter and a software package. It would be used with an existing computer system. Users would install the counter on their computer, using either a supplied interface card or a standard interface like RS-232. They would then install the software. A Loran-C receiver would be optional. If included, it could fit inside the same box as the counter. If not included, the system could still make oscillator-to-oscillator comparisons.

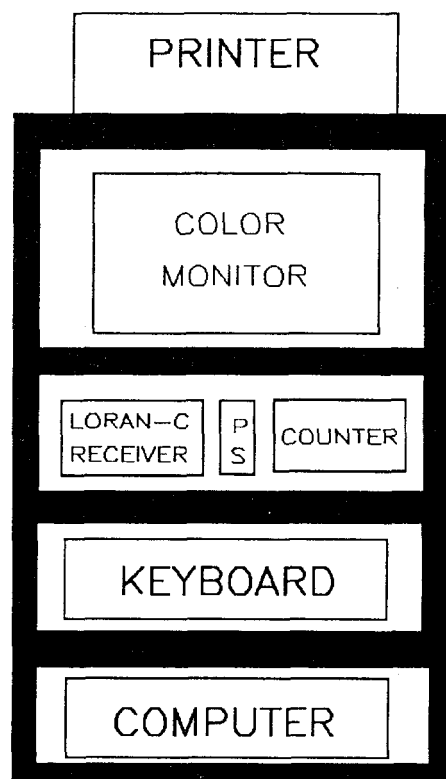


Figure 6 - Layout of existing system

The advantages of system 1 are its neatness and compactness. System 1 is self contained, and requires no installation. However, system 2 is less expensive, since a calibration lab could install it on a PC system that they probably already have. It is more capable, since it has the advantage of using peripherals like big-screen color monitors, hard disks, and laser printers. It is also easier to update. More advanced systems could be designed by writing software for a graphics environment like Microsoft WindowsTM*. And since the hardware is "portable", the system could be moved to other computer platforms by rehosting the software.

CONCLUSION

This paper examined the design philosophy of an automated calibration system, and showed how this type of system benefits the user. NIST hopes to further the cause of laboratory automation by designing easier-to-use and more powerful frequency calibration systems in the years to come.

* Identification of commercial companies in this paper does not imply recommendation or endorsement by the National Institute of Standards and Technology, nor does it imply that any identified entry is the only or the best available for the the purpose.

This paper is a contribution of the the U. S. Government and is not subject to copyright.